

Achievements in Answer Set Programming (Preliminary Report)

Vladimir Lifschitz

University of Texas, Austin, Texas, USA
vl@cs.utexas.edu

Abstract. This paper describes an approach to the methodology of answer set programming (ASP) that can facilitate the design of encodings that are easy to understand and provably correct. Under this approach, after appending a rule or a small group of rules to the emerging program we include a comment that states what has been “achieved” so far. This strategy allows us to set out our understanding of the design of the program by describing the roles of small parts of the program in a mathematically precise way.

1 Introduction

This paper describes an approach to the methodology of answer set programming [13,14] that can facilitate the design of encodings that are easy to understand and provably correct. Under this approach, after appending a rule or a small group of rules to the emerging program, the programmer would include a comment that states what has been “achieved” so far, in a certain precise sense.

Consider, for instance, the program in Fig. 1. Its first rule, viewed as a one-

```
% Program 8Queens

row(1..8).
col(1..8).
8 { queen(I,J) : col(I), row(J) } 8.
:- queen(I,J), queen(I,JJ), J != JJ.
:- queen(I,J), queen(II,J), I != II.
:- queen(I,J), queen(II,JJ), (I,J) != (II,JJ), |I-II|=|J-JJ|.
```

Fig. 1. 8 queens program, adapted from [7, Sect. 3.2].

rule program, has a unique stable model S , which satisfies the following condition:

$$\text{A ground atom of the form } row(i) \text{ belongs to } S \text{ iff } i \in \{1, \dots, 8\}. \quad (1)$$

Condition (1) holds also if S is the stable model of the first two rules of the program in Fig. 1. And it holds if S is any stable model of the first three rules, and so on, for all 6 “prefixes” (initial segments) of the program. This is what we mean by achievement: once the programmer declares that a property “has been achieved,” he is committed to maintaining this property of stable models until the program is completed.

After writing the second rule, the programmer can claim that something else has been achieved:

$$\text{A ground atom of the form } col(j) \text{ belongs to } S \text{ iff } j \in \{1, \dots, 8\}. \quad (2)$$

This condition holds if S is a stable model of any prefix of the program that includes the first two rules.

Additional properties achieved by adding the third rule can be expressed as follows:

$$\begin{aligned} \text{Set } S \text{ contains exactly 8 ground atoms of the form } queen(i, j). \\ \text{For each of these atoms, } i, j \in \{1, \dots, 8\}. \end{aligned} \quad (3)$$

If a program is written in this manner then every achievement documented in the process of writing it describes a property shared by all stable models of the entire program. We conjecture that under some conditions this list of achievements can serve as the skeleton of a proof of its correctness, in the spirit of Edsger Dijkstra’s advice:

... one should not first make the program and then prove its correctness, because then the requirement of providing the proof would only increase the poor programmer’s burden. On the contrary: the programmer should let correctness proof and program grow hand in hand [5].

Recording important achievements in the process of writing an ASP program may be similar to recording important loop invariants in procedural programming: it does not ensure the correctness of the program but helps the programmer move towards the goal of proving correctness.

2 Programs, Prefixes, and Achievements

In this paper, by an (ordered) program we understand a list of rules R_1, \dots, R_n ($n \geq 1$) in the input language of an answer set solver, such as CLINGO [9] or DLV [6]. The order of rules is supposed to reflect the order in which the programmer writes them in the process of creating the program. It does not affect the semantics of the program, but it is essential for understanding the process of programming.

We restrict attention to programs without classical negation. (This limitation is discussed in the conclusion.) Stable models of a program without classical

negation are sets of ground atoms that contain no arithmetic operations, intervals, or pools [9, Sect. 3.1.7, 3.1.9, 3.1.10]. Such ground atoms will be called *precomputed*.¹ An *interpretation* is a set of precomputed atoms.

The k -th prefix of a program R_1, \dots, R_n , where $1 \leq k \leq n$, is the program R_1, \dots, R_k . We will express that a program Γ is a prefix of a program Π by writing $\Gamma \leq \Pi$. The relation \leq is a total order on the set of prefixes of a program.

In some programs, constants are used as placeholders for values provided by the user [9, Sect. 3.1.15]. For example, the constant n is used as a placeholder for an arbitrary positive integer in the generalization of the 8 queens program shown in Fig. 2. We assume here that Π is a program that does not contain

```
% Program NQueens

row(1..n).
col(1..n).
n { queen(I,J) : col(I), row(J) } n.
:- queen(I,J), queen(I,JJ), J != JJ.
:- queen(I,J), queen(II,J), I != II.
:- queen(I,J), queen(II,JJ), (I,J) != (II,JJ), |I-II|=|J-JJ|.
```

Fig. 2. Generalization: N queens program.

such placeholders, so that the stable models of Π are completely determined by its rules, without any additional input from the user. An extension to programs with input is discussed in the next section.

An *achievement* of a prefix Γ of Π is a property of sets of interpretations that holds for all stable models of all programs Δ such that $\Gamma \leq \Delta \leq \Pi$.

For example, (1) is an achievement of the first prefix of program *8Queens* (Fig. 1); (2) is an achievement of its second prefix; and (3) is an achievement of its third prefix. Conditions (1) and (2), and the conjunction of conditions (1)–(3), are achievements of the third prefix as well. Any condition that holds for all sets of interpretations is trivially an achievement of any prefix of any program.

The following three conditions are achievements of the last three prefixes of *8Queens*:

$$\begin{aligned} &\text{Each column of the } 8 \times 8 \text{ chessboard includes at most one square } (i, j) \\ &\text{such that the atom } \textit{queen}(i, j) \text{ belongs to } S. \end{aligned} \quad (4)$$

¹ This terminology follows Gebser et al. [7, Sect. 2.1], where “precomputed terms” are defined. Calimeri et al. [3, Sect. 2.1] talk about elements of the “Herbrand universe” of a program in the same sense.

Each row of the 8×8 chessboard includes at most one square (i, j) such that the atom $queen(i, j)$ belongs to S . (5)

Each diagonal of the 8×8 chessboard includes at most one square (i, j) such that the atom $queen(i, j)$ belongs to S . (6)

Thus every stable model S of $8Queens$ satisfies all conditions (1)–(6).

3 Programs with Input

The value of a placeholder, such as constant n in Fig. 2, is one kind of input that an answer set solver may expect in addition to the rules of the program. A definition of an “extensional predicate” occurring in the bodies of rules is another kind. For example, the program *Hamiltonian* (Fig. 3) needs to be supplemented

```
% Program Hamiltonian

1 {in(X,Y) : edge(X,Y) } 1 :- vertex(X).
1 {in(X,Y) : edge(X,Y) } 1 :- vertex(Y).
reached(X) :- in(v0,X).
reached(Y) :- reached(X), in(X,Y).
:- not reached(X), vertex(X).
```

Fig. 3. Hamiltonian cycle program, adapted from [7, Sect. 3.3].

by definitions of

- the predicate $vertex/1$, representing the set of vertices of a finite directed graph G ,
- the predicate $edge/2$, representing the set of edges of G ,
- the placeholder $v0$, which is a vertex of G .

If \mathbf{i} is a valid input for a program Π then we can talk about models of Π that are stable for input \mathbf{i} , or “ \mathbf{i} -stable.”²

For a program with input, an achievement is defined as a relation between valid inputs and sets of interpretations. Such a relation R will be called an achievement of a prefix Γ of Π if $R(\mathbf{i}, S)$ holds whenever S is an \mathbf{i} -stable model of a program Δ such that $\Gamma \leq \Delta \leq \Pi$.

For example, the sentence

$$\text{A ground atom of the form } row(i) \text{ belongs to } S \text{ iff } i \in \{1, \dots, n\} \quad (7)$$

² Programs with input are similar to lp-functions in the sense of Gelfond and Przymusińska [11].

expresses a relation between n and S that is an achievement of the first prefix of *NQueens*. It is obtained from condition (1) by replacing 8 with n , and achievements of the other prefixes of that program can be obtained in a similar way from conditions (2)–(6).

The following two conditions are achievements of the first two prefixes of *Hamiltonian*:

Every pair (x, y) such that the atom $in(x, y)$ belongs to S is an edge of G ; for every vertex x of G there is a unique y such that the atom $in(x, y)$ belongs to S . (8)

For every vertex y of G there is a unique x such that the atom $in(x, y)$ belongs to S . (9)

Nothing interesting has been achieved by adding the third rule, but the following condition is an achievement of the fourth prefix of the program:

The set of symbols x such that the atom $reached(x)$ belongs to S consists of the vertices x for which there exists a walk v_0, \dots, v_n such that $n \geq 1$, $v_0 = v_0$, $v_n = x$, and every atom of the form $in(v_i, v_{i+1})$ belongs to S . (10)

Finally, here is an achievement of the entire program:

For every vertex x of G , the atom $reached(x)$ belongs to S . (11)

4 Records of Achievement

A *record of achievement* for a program Π is a function that maps some (possibly all) prefixes of Π , including Π itself, to their achievements. For instance, the function that maps the prefixes of *8Queens* to conditions (1)–(6) is a record of achievement, as well as the function that maps the first, second, fourth and fifth prefixes of *Hamiltonian* to conditions (8)–(11). Every program has a trivial record of achievement that maps all its prefixes to the identically true condition.

A record of achievement a can be represented by including in the program, after each prefix Γ in the domain of a , a comment that describes the condition $a(\Gamma)$. In Fig. 4 we use a convention that helps us describe these conditions concisely: in the comment

```
achieved: row/1 = {1,...,n}
```

we understand *row/1* as shorthand for “the set of precomputed terms i such that the atom $row(i)$ belongs to S .” In view of this convention, the comment above is a concise reformulation of condition (7). In the other comments, *col/1* and *queen/2* are understood in a similar way.

Figure 4 includes also a comment that tells us which inputs for the program are considered valid.

Figure 5 uses another useful convention. In the comment

```

% Program NQueens

% input: positive integer n (size of the board).

% A square is represented as a pair, column number and row number, both
% from the set {1,...,n}.

row(1..n).
% achieved: row/1 = {1,...,n}.

col(1..n).
% achieved: col/1 = {1,...,n}.

n { queen(I,J) : col(I), row(J) } n.
% achieved: Set queen/2 consists of n squares.

:- queen(I,J), queen(I,JJ), J != JJ.
% achieved: Each column includes at most one square from queen/2.

:- queen(I,J), queen(II,J), I != II.
% achieved: Each row includes at most one square from queen/2.

:- queen(I,J), queen(II,JJ), (I,J) != (II,JJ), |I-II|=|J-JJ|.
% achieved: Each diagonal includes at most one square from queen/2.

```

Fig. 4. The N queens program with a record of achievement.

```

% Program Hamiltonian

% input: the set vertex/1 of vertices of a finite directed graph G.
% input: the set edge/2 of edges of G.
% input: a vertex v0 of G.

1 {in(X,Y) : edge(X,Y) } 1 :- vertex(X).
% achieved: Set in/2 is a subset of edge/2; for every vertex X of G
%           there is a unique Y such that in(X,Y).

1 {in(X,Y) : edge(X,Y) } 1 :- vertex(Y).
% achieved: For every vertex Y of G there is a unique X such that in(X,Y).

reached(X) :- in(v0,X).
reached(Y) :- reached(X), in(X,Y).
% achieved: Set reached/1 consists of the vertices that are reachable
%           from v0 by a path of non-zero length in the subgraph of G
%           with the set of edges in/2.

:- not reached(X), vertex(X).
% achieved: reached/1 = vertex/1.

```

Fig. 5. The Hamiltonian cycle program with a record of achievement.

achieved: For every vertex Y of G there is a unique X such that $\text{in}(X,Y)$

we understand X and Y as metavariables for precomputed terms, and this comment is interpreted as follows: “for every vertex y of G , there exists a unique precomputed term x such that the atom $\text{in}(x,y)$ belongs to S .” Thus the comment above is a reformulation of condition (9).

5 Completeness

If a is a record of achievement for a program Π , and Δ is a prefix of Π , then by $a^*(\Delta)$ we denote the conjunction of conditions $a(\Gamma)$ for all prefixes Γ of Δ that belong to the domain of a . It is clear that condition $a^*(\Delta)$ is an achievement of Δ . In particular, all stable models of Π satisfy condition $a^*(\Pi)$. The converse is, generally, not true—it is possible that an interpretation satisfying $a^*(\Pi)$ is not a stable model of Π .

The following notation and terminology will be used to discuss this issue. For any prefix Γ of a program Π , by $\text{Preds}(\Gamma)$ we denote the set consisting of the predicates occurring in Γ and the input predicates of Π . For instance, if Γ is the first prefix of *Hamiltonian* then $\text{Preds}(\Pi)$ consists of the predicates

$$\text{vertex}/1, \text{edge}/2, \text{in}/2. \quad (12)$$

The expression $\text{Preds}(I)$, where I is an interpretation, stands for the set of predicates occurring in I . About a record of achievement a for a program Π we will say that it is *complete* if for every Γ in the domain of a , all interpretations I satisfying $a^*(\Gamma)$ for which $\text{Preds}(I) \subseteq \text{Preds}(\Gamma)$ are stable models of Γ .

The records of achievement shown in Fig. 4 and 5 are complete. The completeness of the latter, for example, entails that

- (a) any interpretation S such that atoms in S contain no predicates other than (12) is a stable model of the first rule of *Hamiltonian* iff it satisfies condition (8);
- (b) such an interpretation is a stable model of the first two rules of *Hamiltonian* iff it satisfies conditions (8) and (9);
- (c) any interpretation S such that atoms in S contain no predicates other than

$$\text{vertex}/1, \text{edge}/2, \text{in}/2, \text{reached}/1$$

is a stable model of the first four rules of *Hamiltonian* iff it satisfies conditions (8)–(10);

- (d) such an interpretation is a stable model of the entire program iff it satisfies conditions (8)–(11).

The completeness property of the record of achievement in Fig. 5 is closely related to the correctness of the program as an encoding of Hamiltonian cycles. From property (d) we can conclude that a set of edges of G is a Hamiltonian cycle iff it has the form

$$\{(x, y) : \text{in}(x, y) \in S\}$$

for some stable model S of *Hamiltonian*.

6 Achievement-Based Answer Set Programming

Records of achievement in Fig. 4 and 5 are not only complete but also detailed, in the sense that they include achievements for almost all prefixes of the programs. The only rule in these programs that is not followed by an achievement comment is the first rule in the recursive definition of *reached*. The role of that rule cannot be properly explained unless we treat it as part of the definition.

Developing an ASP program along with a complete and detailed record of achievement can be called “achievement-based” answer set programming. This strategy allows us to set out our understanding of the design of the program by describing the roles of individual rules, or small groups of rules, in a mathematically precise way.

To further illustrate this idea, we present in Fig. 6–8 three “real life” ASP programs accompanied by complete, detailed records of achievement. In Fig. 7, the rule

$$:- \text{hb}(\text{N}, \text{X}, \text{Y}), \text{hb}(\text{N}, \text{Y}, \text{Z}), \text{not } \text{hb}(\text{N}, \text{X}, \text{Z}).$$

replaces the pair of rules

```

% Program OBT

% input: positive integer k.

leaf(0..k).
% achieved: leaf/1 = {0,...,k}.

vertex(0..2*k).
% achieved: vertex/1 = {0,...,2k}.

internal(X) :- vertex(X), not leaf(X).
% achieved: internal/1 = {k+1,...,2k}.

2 {edge(X,Y) : vertex(Y), X>Y} 2 :- internal(X).
% Let G be the digraph with the vertices vertex/1 and the edges edge/2.
% achieved: for every edge (X,Y) of G, X>Y; the out-degree of a vertex
%           X in G is 2 if internal(X), and 0 if leaf(X).

reachable(X,Y) :- edge(X,Y).
reachable(X,Y) :- edge(X,Z), reachable(Z,Y).
% achieved: reachable(X,Y) iff Y is reachable from X in G by a path of
%           non-zero length.

:- vertex(X), X!=2*k, not reachable(2*k,X).
% achieved: every vertex of G other than 2k is reachable from 2k by a
%           path of non-zero length.

:- reachable(X,X), vertex(X).
% achieved: G is acyclic.

max_child(X,Y) :- edge(X,Y), edge(X,Y1), Y > Y1.
% achieved: max_child(X,Y) iff Y is the largest child of X in G.

Y<Y1 :- max_child(X,Y), max_child(X1,Y1), Y>Y1, X<X1.
% achieved: for any vertices X, X1 of G such that X<X1, the largest
%           child of X is smaller than the largest child of X1.

```

Fig. 6. Encoding of ordered binary trees, adapted from [2, Sect. 1], with a complete, detailed record of achievement. An ordered binary tree is a rooted binary tree with the leaves $0, \dots, k$ and internal vertices $k+1, \dots, 2k$ such that (i) every internal vertex is greater than its children, and (ii) for any two internal vertices x and x_1 , $x > x_1$ iff the maximum of the children of x is greater than the maximum of the children of x_1 .

```

% Program SCA

% input: the number s of symbols 1,...,s.
% input: the number n of rows 1,...,n.

sym(1..s).
% achieved: sym/1 = {1,...,s}.

row(1..n).
% achieved: row/1 = {1,...,n}.

1 {hb(N,X,Y); hb(N,Y,X)} 1 :- row(N), sym(X), sym(Y), X!=Y.
% For every row N, let hb_N be the binary relation on sym/1 defined by
% the condition: X hb_N Y iff hb(N,X,Y).
% achieved: each relation hb_N is irreflexive; each pair of distinct
%           symbols satisfies either X hb_N Y or Y hb_N X.

:- hb(N,X,Y), hb(N,Y,Z), not hb(N,X,Z).
% achieved: each relation hb_N is transitive.

covered(X,Y,Z) :- hb(N,X,Y), hb(N,Y,Z).
% For every row N and every symbol X, by M_{N,X} we denote the symbol
% that is the X-th smallest w.r.t. hb_N.
% achieved: for any symbols X, Y, Z, covered(X,Y,Z) iff, for some row N,
%           (X,Y,Z) is a subsequence of (M_{N,1},...,M_{N,s}).

:- not covered(X,Y,Z), sym(X), sym(Y), sym(Z), X!=Y, Y!=Z, X!=Z.
% achieved: covered(X,Y,Z) for any pairwise distinct symbols X, Y, Z.

```

Fig. 7. Encoding of sequence covering arrays of strength 3, adapted from [1, Fig. 1], with a complete, detailed record of achievement. A sequence covering array of strength t is an array of permutations of symbols such that every ordering of any t symbols appears as a subsequence of at least one row.

```

% Program Borda

% input: the number m of candidates 1,...,m in an election E.
% input: the set p/3 of triples (P,Pos,C) such that, for a fixed ordering
%       pr_1,...,pr_l of the distinct preference relations in the profile
%       of E, candidate C is at position Pos in relation pr_P.
% input: the set votecount/2 of pairs (P,VC) such that relation pr_P
%       occurs VC times in the profile of E.

candidate(1..m).
% achieved: candidate/1 = {1,...,m}.

posScore(P,C,X*VC) :- p(P,Pos,C), X=m-Pos, votecount(P,VC).
% achieved: posScore(P,C,S) iff the voters who chose relation pr_P in
%           election E contributed S points to candidate C under the
%           Borda rule.

score(C,N) :- candidate(C), N=#sum{S:posScore(_,C,S)}.
% achieved: score(C,N) iff candidate C earned N points in election E
%           under the Borda rule.

winner(C) :- score(C,M), M=#max{S:score(_,S)}.
% achieved: winner(C) iff the number of points earned by candidate C in
%           election E is maximal among all candidates.

```

Fig. 8. Encoding of the Borda voting rule, adapted from [4, Encoding 1], with a complete, detailed record of achievement. Each voter ranks the list of candidates in order of preference. The candidate ranked last gets zero points; next to last gets one point, and so on. The candidate with the most points is the winner.

$$\begin{aligned} \text{hb}(N, X, Z) &:- \text{hb}(N, X, Y), \text{hb}(N, Y, Z). \\ &:- \text{hb}(N, X, X). \end{aligned}$$

from the original paper by Brain et al. [1]. The problem with the original version is that the first rule of the pair may temporarily destroy the irreflexivity of the relation of hb_N that was true at the previous step; this property is restored by the second rule. That is not in the spirit of the achievement-based approach, which emphasizes the gradual accumulation of properties that we would like to see in the complete program.

7 Conclusion

As we are adding rules to an emerging ASP program, we deal at every step with a single executable piece of code, unlike the non-executable pseudo-code formed in the process of stepwise refinement of a procedural program, and unlike a collection of executable subroutines formed in the process of bottom-up design. The idea of looking at a prefix as if it was a complete program, thinking about its stable models, and relating them to the stable models of the final product is at the root of achievement-based ASP.

When a program is developed in accordance with this approach, it begins with comments that precisely describe its valid inputs. Then, after every rule or small group of rules, we include a comment describing what has been achieved. Collectively these comments represent a complete record of achievement.

According to Gebser et al. [7, Sect. 3.2],

[t]he basic approach to writing encodings in ASP follows a *generate-and-test* methodology, also referred to as *guess-and-check*. . . A “generating” part is meant to non-deterministically provide solution candidates, while a “testing” part eliminates candidates violating some requirements. . . Both parts are usually amended by “defining” parts providing auxiliary concepts.

Most programs discussed in this paper are designed in accordance with this basic approach.³ In particular, the two rules with *reached* in the head (Fig. 3) define an auxiliary concept, as well as two rules with *reachable* in the head in Fig. 6 and the rule with *covered* in the head in Fig. 7. When a program includes a definition consisting of more than one rule then it has a prefix that covers only a part of the definition, and such a prefix is likely to achieve nothing of interest, as in the case of *Hamiltonian* and *OBT*.

³ The program in Fig. 8 is an exception—it has no generating part and no testing part. Also, it is not clear whether the designers of *Hamiltonian* (Fig. 3) intended the second rule for the generating part or for the testing part. (The second rule is syntactically similar to the first, which is definitely a generate rule. On the other hand, adding the second rule does not really generate new solution candidates; it eliminates some of the candidates generated earlier.)

Our advice—as you are adding rules to your program, keep track of what has been achieved—differs from the “generate-and-test” advice in that it refers to mathematical properties of stable models, and not to programmer’s intentions.

The programs discussed in this paper do not use classical negation. In the presence of classical negation, answer sets consist of “precomputed classical literals”—precomputed atoms and classical negations of such atoms. Extending the definition of a complete record of achievement to such programs is straightforward. On the other hand, many programs with classical negation contain defaults [10, Chap. 5], such as the closed world assumption and the common-sense law of inertia, and the achievement-based approach may be not so useful in application to programs containing defaults. A default does not “achieve” anything in the technical sense of Sect. 2.

In this preliminary report, claims about records of achievement and their completeness are stated without proof. We would like to develop theoretical tools for proving such claims. Proofs of this kind will rely on precise descriptions of the semantics of underlying programming languages, such as the definitions of the core language of DLV [12, Sect. 2.2], ASP Core [3, Sect. 2], or Abstract Gringo [8, Sect. 4]. (These definitions need to be extended to programs with input.) Such proofs will rely also on precise definitions of the terminology used in achievement comments (“diagonals of the chessboard,” “the Borda rule”).

If a program has a precise specification then we would like its record of achievement to be closely related to that specification, so that the correctness of the program will easily follow from the completeness of the record. In such cases, the record of achievement may suggest a way to organize a proof of correctness: go over the achievements in its domain one by one, and prove that each of them correctly describes the stable models of the corresponding prefix of the program. Once we have arrived at the description of the stable models of the entire program, we will complete the proof of correctness by relating that description to the given specification.

Acknowledgements

Thanks to Michael Gelfond, Amelia Harrison, Yuliya Lierler, Julian Michael, and Liangkun Zhao for comments on drafts of this paper. Conversations and exchanges of email messages with Mark Denecker, Esra Erdem, Martin Gebser, Roland Kaminski, Johannes Oetsch, Dhananjay Raju, and Mirek Truszczynski, helped the author develop a better understanding of the methodology of answer set programming. This research was partially supported by the National Science Foundation under Grant IIS-1422455.

References

1. Brain, M., Erdem, E., Inoue, K., Oetsch, J., Pührer, J., Tompits, H., Yilmaz, C.: Event-sequence testing using answer-set programming. *International Journal on Advances in Software* 5, 237–251 (2012)

2. Brooks, D.R., Erdem, E., Erdoğan, S.T., Minett, J.W., Ringe, D.: Inferring phylogenetic trees using answer set programming. *Journal of Automated Reasoning* 39, 471–511 (2007)
3. Calimeri, F., Faber, W., Gebser, M., Ianni, G., Kaminski, R., Krennwallner, T., Leone, N., Ricca, F., Schaub, T.: ASP-Core-2: Input language format. Available at <https://www.mat.unical.it/aspcomp2013/files/ASP-CORE-2.03b.pdf> (2012)
4. Charwat, G., Pfandler, A.: Democratix: A declarative approach to winner determination. In: *Proc. of the 4th International Conference on Algorithmic Decision Theory (ADT 2015)* (2015), accepted for publication.
5. Dijkstra, E.W.: The humble programmer. *Communications of the ACM* 15, 859–866 (1972)
6. Eiter, T., Leone, N., Mateis, C., Pfeifer, G., Scarcello, F.: The KR system DLV: Progress report, comparisons and benchmarks. In: Cohn, A., Schubert, L., Shapiro, S. (eds.) *Proceedings of International Conference on Principles of Knowledge Representation and Reasoning (KR)*. pp. 406–417 (1998)
7. Gebser, M., Kaminski, R., Kaufmann, B., Schaub, T.: *Answer Set Solving in Practice*. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, Morgan and Claypool Publishers (2012)
8. Gebser, M., Harrison, A., Kaminski, R., Lifschitz, V., Schaub, T.: Abstract Gringo. *Theory and Practice of Logic Programming* 15, 449–463 (2015)
9. Gebser, M., Kaminski, R., Kaufmann, B., Lindauer, M., Ostrowski, M., Romero, J., Schaub, T., Thiele, S.: *Potassco User Guide*, version 2.0 (2015), available at <http://potassco.sourceforge.net>
10. Gelfond, M., Kahl, Y.: *Knowledge Representation, Reasoning, and the Design of Intelligent Agents: The Answer-Set Programming Approach*. Cambridge University Press (2014)
11. Gelfond, M., Przymusinska, H.: Towards a theory of elaboration tolerance: Logic programming approach. *International Journal of Software Engineering and Knowledge Engineering* 6(1), 89–112 (1996)
12. Leone, N., Pfeifer, G., Faber, W., Eiter, T., Gottlob, G., Perri, S., Scarcello, F.: The DLV system for knowledge representation and reasoning. *ACM Transactions on Computational Logic* 7(3), 499–562 (2006)
13. Marek, V., Truszczyński, M.: Stable models and an alternative logic programming paradigm. In: *The Logic Programming Paradigm: a 25-Year Perspective*, pp. 375–398. Springer Verlag (1999)
14. Niemelä, I.: Logic programs with stable model semantics as a constraint programming paradigm. *Annals of Mathematics and Artificial Intelligence* 25, 241–273 (1999)